

CLAIMS

What is claimed is:

1. A data processing system that includes a network comprised of a set of connected devices, individual ones of said connected devices of said set of connected devices comprising a data processor that executes a program for connecting to and disconnecting from said network and for maintaining a first list descriptive of other connected devices and a second list descriptive of at least some user-defined services published by individual ones of publisher connected devices that are members of a first sub-set of the connected devices, individual ones of said publisher connected devices maintaining a third list descriptive of an individual one or individual ones of registered service subscriber connected devices that are members of a second sub-set of connected devices, said publisher connected devices providing a service output to registered service subscriber connected devices upon an occurrence of at least one predetermined triggering event, wherein a given one of said connected devices may be a member of only said first sub-set of connected devices, or may be a member of only said second sub-set of connected devices, or may be a member of both said first sub-set of connected devices and said second sub-set of connected devices.
2. A data processing system as in claim 1, wherein individual ones of said publisher connected devices of said first sub-set of connected devices have at least one back-up connected device to publish the service upon the unavailability of said publisher connected device, said at least one back-up connected device storing data and procedures required to publish the service should the publisher connected device become unavailable.
3. A data processing system as in claim 1, wherein said triggering event is comprised of an occurrence of a push trigger defined by said publisher connected device.
4. A data processing system as in claim 1, wherein said triggering event is comprised of an occurrence of a pull trigger defined by said subscriber connected device.

5. A data processing system as in claim 1, wherein said triggering event is comprised of an occurrence of a push trigger defined by said publisher connected device and by an occurrence of a pull trigger defined by said subscriber connected device.
6. A data processing system as in claim 1, wherein a content of said first list is changed automatically as connected devices connect to and disconnect from said network.
7. A data processing system as in claim 1, wherein at least one of said published services is based on the operation of a spreadsheet engine operated by one of said publisher connected devices.
8. A data processing system as in claim 1, wherein said network is comprised of a Local Area Network (LAN), and wherein said connected devices self-register on the LAN when connecting to the LAN using Mailslots.
9. A data processing system as in claim 1, wherein said network is comprised of a Wide Area Network (WAN), and wherein said connected devices self-register on the WAN when connecting to the WAN using Named Pipes.
10. A data processing system as in claim 1, wherein at least one of said published services retrieves data in response to a database query originated by a subscriber connected device, wherein the data is retrieved from a database where the publisher connected device has access rights.
11. A data processing system as in claim 1, wherein at least one of said published services is offered by a first publisher connected device that subscribes to at least one second publisher connected device, wherein a report generated by said first publisher connected device comprises a portion of a report received from the second publisher connected device.
12. A data processing system as in claim 1, wherein at least one of said published services is offered by a first publisher connected device that subscribes to at least one second publisher connected device, wherein a report generated by said first publisher connected device comprises a processed portion of a report received from the second publisher connected device.

13. A data processing system as in claim 1, wherein set of connected devices communication over said network using TCP/IP.
14. A data processing system as in claim 1, wherein at one of said set of connected devices is coupled to said network through a wireless link.
15. A data processing system as in claim 1, wherein at least one of said connected devices comprises a measuring machine that generates data and publishes the data in a report service to which at least one publisher connected device subscribes, said subscribing publisher connected device offering a service that reports all or a portion of the data generated by said measuring machine.
16. A data processing system as in claim 1, wherein at least one of said connected devices comprises a measuring machine that generates data and publishes the data in a report service comprising a spreadsheet to which at least one publisher connected device subscribes, said subscribing publisher connected device executing a spreadsheet engine that treats at least some of said data as spreadsheet variables and offers a service that reports a result of a processing of said spreadsheet variables by said spreadsheet engine.
17. A data processing system as in claim 16, wherein a triggering event comprises a change in state in at least one of said spreadsheet variables.
18. A data processing system as in claim 1, wherein a database is stored in a distributed manner by a plurality of said connected devices, and wherein subscribers to said database issue queries that are responded to in a distributed manner by said plurality of connected devices.
19. A data processing system as in claim 1, wherein a publisher connected device maintains a data warehouse by subscribing to at least one publisher connected device that is coupled to a database, and responds to a query from a connected device for data stored in said data warehouse.
20. A data processing system as in claim 1, wherein at least one database is accessible to a publisher connected device to which a subscriber connected device issues queries, said publisher connected

device issuing a database query on behalf of the subscriber connected device and returns a result of the query to the subscriber connected device.

21. A data processing system as in claim 1, wherein a database is maintained by at least one publisher connected device, said database storing data related to a project, and wherein a subscriber connected device has access to at least a portion of said project-related data regardless of its storage location, so long as said publisher connected device has access to the project-related data.

22. A data processing system as in claim 1, wherein a service offered by a publisher connected device migrates to at least one backup connected device that forms a part of a hierarchy of backup connected devices.

23. A data processing system as in claim 1, wherein one of said connected devices of said hierarchy of backup connected devices is designated as a persistent backup connected device.

24. A data processing system as in claim 22, wherein a connected device's location within the hierarchy of backup connected devices is determined based on a plurality of criteria, including the availability and available processing power of each connected device being considered for placement into the hierarchy of backup connected devices.

25. A data processing system as in claim 24, wherein both long term and short term scores are derived using said criteria.

26. A data processing system as in claim 23, wherein said persistent backup connected device is comprised of a mainframe computer.

27. A data processing system as in claim 1, wherein a database is maintained by at least one publisher connected device, said database storing data related to a project, and wherein a subscriber connected device has access to at least a portion of said project-related data, including project-related data stored in private folders, so long as said publisher connected device has access to the project-related data.

28. A data processing system as in claim 1, wherein a publisher defines a service through the use of a scripting engine.
29. A data processing system as in claim 1, wherein a reporting of services to subscriber connected devices is conducted through a user interface that comprises a browser in the subscriber connected device that views web pages provided by a server of the publisher connected device.
30. A data processing system as in claim 1, wherein said system is COM-compliant.
31. A data processing system as in claim 1, wherein at least one service is provisioned, at least in part, by interoperability with a COM-compliant application program.
32. A data processing system as in claim 1, wherein individual ones of said connected devices comprise middleware software and function as middleware servers.
33. A data processing system as in claim 1, wherein said set of connected devices are logically organized into a plurality of clusters each having a top level connected device (TLCD), and wherein said plurality of TLCDs are coupled together in a ring for propagating system administration information between said plurality of clusters.
34. A data processing system as in claim 33, wherein said system administration information is used at least in part for changing a content of said first list as connected devices connect to and disconnect from said network.
35. A data processing system as in claim 1, wherein at least one of said published services is a data retrieval service wherein a connected device having privileges to access a database employs a query tool to compose a query, wherein said query is placed at the disposal of a subscriber connected device who does not have privileges to access the database.
36. A data processing system as in claim 1, wherein at least one of said published services is a data

retrieval service wherein a connected device having privileges to access a database is responsive to a query composed by a subscriber connected device, who does not have privileges to access the database, to access the database using the query and to return the results of the query to the subscriber connected device.

37. A data processing system as in claim 1, wherein a service offered by a publisher connected device migrates to at least one backup connected device that forms a part of a hierarchy of backup connected devices, wherein a determination of the hierarchy of appropriate backup connected devices is based at least in part on a score derived from long term loads, and wherein balancing of the load on backup connected devices is accomplished at least in part by limiting the number of services that may become active based on a score derived from short term loads.

38. A data processing system as in claim 33, wherein a service offered by a publisher connected device migrates to at least one backup connected device, wherein the backup connected device is constrained to reside within the cluster of the publisher connected device.

39. A data processing system as in claim 1, wherein a reporting service is implemented by a fanout mechanism so as to reduce the load on a publisher connected device.

40. A data processing system as in claim 1, wherein said set of connected devices are logically organized into a plurality of clusters each having a top level connected device (TLCD), and wherein a connected device is enabled to locate a service with certain attributes by a drill down procedure through information about a service provided by a publisher connected device.

41. A data processing system as in claim 1, wherein said set of connected devices are logically organized into a plurality of clusters each having a top level connected device (TLCD), and wherein a connected device is enabled to locate a service with certain attributes by issuing a query that is passed all TLCDs, and TLCDs pass the query to all cluster members connected devices, and those cluster member connected devices that return an output do so to the querier connected device in a manner that is unmediated by said TLCDs.

03049051 "050444

42. A data processing system as in claim 1, wherein in response to a connected device applying to a publisher connected device for a subscription, said publisher connected device one of selectively grants the subscription or grants the subscription automatically.

43. A data processing system as in claim 1, wherein at least one publisher connected device provides a service that is responsive to commands native to at least one of an Enterprise Resource Planning (ERP) system, a Material Requirements Planning (MRP) system, or a Manufacturing Execution System (MES).

44. A data processing system as in claim 33, wherein said clusters are constructed in parallel as connected devices request connection to the system so as to add connected devices in turn to individual ones of N clusters.

45. A data processing system as in claim 33, wherein said clusters are constructed sequentially so as to add connected devices to one cluster until the cluster reaches a predetermined size, and to then add the next connected device to a next cluster.

46. A data processing system as in claim 23, wherein said persistent backup connected device is comprised of a central server.

47. A data processing system as in claim 1, wherein a reporting of services to subscriber connected devices is conducted by a web server of the publisher connected device that hosts a report in the form of a web page.

48. A data processing system that includes a network comprised of a set of connected devices, individual ones of said connected devices of said set of connected devices comprising a data processor that executes a program for connecting to and disconnecting from said network and for maintaining a first list descriptive of other connected devices and a second list descriptive of at least some user-defined services published by individual ones of publisher connected devices that are members of a first sub-set of the connected devices, individual ones of said publisher connected devices maintaining a third list descriptive of an individual one or individual ones of registered

service subscriber connected devices that are members of a second sub-set of connected devices, wherein at least one publisher connected device offers a data retrieval service from at least one database, and wherein at least one service subscriber connected device is registered to receive said data retrieval service.

49. A data processing system as in claim 48, wherein said at least one publisher connected device provides retrieved data to a registered service subscriber connected device upon an occurrence of at least one predetermined triggering event.

50. A data processing system as in claim 48, wherein said data retrieval service automatically moves to at least one backup connected device at least upon an occurrence of the publisher connected device becoming unable to provide the data retrieval service.

51. A data processing system as in claim 50, said at least one backup connected device is one of a hierarchy of backup connected devices, one of which is designated as a persistent backup connected device.

52. A data processing system as in claim 48, wherein said publisher connected device makes a database query on behalf of the subscriber connected device who does not have privileges to access the database.

53. A data processing system as in claim 48, wherein said set of connected devices are logically organized into a plurality of clusters each having a top level connected device (TLCD), and wherein said plurality of TLCDs are coupled together in a ring for propagating system administration information between said plurality of clusters.

54. A data processing system as in claim 48, wherein said publisher connected device makes a database query on behalf of the subscriber connected device, where the subscriber connected device is coupled to said network through a wireless link.

55. A data processing system as in claim 48, wherein said database stores data related to a project,

and wherein said subscriber connected device has access to at least a portion of said project-related data regardless of its storage location, so long as said publisher connected device has access to the project-related data.

56. A data processing system that includes a network comprised of a set of connected devices, individual ones of said connected devices of said set of connected devices comprising a data processor that executes a program for connecting to and disconnecting from said network and for maintaining a first list descriptive of other connected devices and a second list descriptive of at least some user-defined services published by individual ones of publisher connected devices that are members of a first sub-set of the connected devices, individual ones of said publisher connected devices maintaining a third list descriptive of an individual one or individual ones of registered service subscriber connected devices that are members of a second sub-set of connected devices, wherein at least one publisher connected device offers a project data management retrieval service providing access to project management data that comprises project-related files and communications, and wherein at least one service subscriber connected device is registered to use said project management data retrieval service and to have access to at least a portion of said project management data regardless of its storage location, so long as said at least one publisher connected device has access to the project-related data.

57. A data processing system as in claim 56, wherein said at least one publisher connected device provides the project management data to a registered service subscriber connected device upon an occurrence of at least one predetermined triggering event.

58. A data processing system as in claim 56, wherein said project data management retrieval service automatically moves to at least one backup connected device at least upon an occurrence of the publisher connected device becoming unable to provide the project data management retrieval service.

59. A data processing system as in claim 58, wherein said at least one backup connected device is one of a hierarchy of backup connected devices, one of which is designated as a persistent backup connected device.

60. A data processing system as in claim 56, wherein said set of connected devices are logically organized into a plurality of clusters each having a top level connected device (TLCD), and wherein said plurality of TLCDs are coupled together in a ring for propagating system administration information between said plurality of clusters.

61. A data processing system as in claim 56, wherein said publisher connected device retrieves project management data for the subscriber connected device, and sends the retrieved project management data to the subscriber connected device through a wireless link.

62. A method for operating a network comprised of a plurality of connected devices, comprising:

in each connected device, maintaining a first list descriptive of other connected devices and a second list descriptive of at least some user-defined services published by individual ones of publisher connected devices;

in each publisher connected device, maintaining a third list descriptive of an individual one or individual ones of service subscriber connected devices that are registered to receive the service published by said publisher connected device upon an occurrence of at least one predetermined triggering event;

logically organizing the plurality of connected devices into a plurality of clusters each having a top level connected device, wherein said plurality of top level connected devices are coupled together in a ring structure for propagating system administration information between said plurality of clusters; and

automatically moving the service published by said publisher connected device to at least one backup connected device at least upon an occurrence of the publisher connected device becoming unable to provide the service.

63. A method as in claim 62, wherein said at least one backup connected device is one of a hierarchy of backup connected devices, and further comprising designating one of said hierarchy

of backup connected devices to be a persistent backup connected device.

64. A method as in claim 62, wherein said publisher connected device offers a project data management retrieval service providing access to project management data that comprises project-related files and communications, and wherein at least one service subscriber connected device is registered to use said project management data retrieval service and to have access to at least a portion of said project management data regardless of its storage location, so long as said at least one publisher connected device has access to the project-related data.

65. A method as in claim 64, wherein said publisher connected device retrieves project management data for the subscriber connected device, and further comprising sending the retrieved project management data to the subscriber connected device through a wireless link.

66. A method as in claim 62, wherein said publisher connected device offers a data retrieval service from at least one database, and wherein at least one service subscriber connected device is registered to receive said data retrieval service.

67. A method as in claim 66, wherein said publisher connected device makes a database query on behalf of the subscriber connected device who does not have privileges to access the database.

68. A method as in claim 66, wherein said publisher connected device makes a database query on behalf of the subscriber connected device, and further comprising sending a result of query to the subscriber connected device through a wireless link.

69. A method as in claim 66, wherein said database stores data related to a project, and wherein said subscriber connected device has access to at least a portion of said project-related data regardless of its storage location, so long as said publisher connected device has access to the project-related data.

70. A method as in claim 62, wherein said publisher connected device offers a data mining service for data stored in at least one database, and wherein at least one service subscriber connected device

is registered to use said data mining service.

71. A method as in claim 62, wherein said publisher connected device offers a data warehouse service for data obtained from at least one database, and wherein at least one service subscriber connected device is registered to use said data warehouse service.

72. A method as in claim 62, wherein said publisher connected device offers a service based on a spreadsheet engine, and wherein at least one service subscriber connected device receives a report upon an occurrence of a change of state of at least one spreadsheet variable.

73. A method as in claim 62, wherein said publisher connected device provides a data retrieval service from at least one database, wherein at least one service subscriber connected device is registered for generating queries to said data retrieval service, and wherein said publisher connected device functions as a proxy to make a database query on behalf of the subscriber connected device.

74. A method as in claim 62, and further comprising accounting for use of a service provided by said publisher connected device.